arXiv:1512.00150v1 [math.ST] 1 Dec 2015

Optimal Estimation and Completion of Matrices with Biclustering Structures

Chao Gao[1], Yu Lu[1], Zongming Ma[2], and Harrison H. Zhou[1]

[1] *Yale University*
[2] *University of Pennsylvania*

Abstract

Biclustering structures in data matrices were first formalized in a seminal paper by John Hartigan [12] where one seeks to cluster cases and variables simultaneously. Such structures are also prevalent in block modeling of networks. In this paper, we develop a unified theory for the estimation and completion of matrices with biclustering structures, where the data is a partially observed and noise contaminated data matrix with a certain biclustering structure. In particular, we show that a constrained least squares estimator achieves minimax rate-optimal performance in several of the most important scenarios. To this end, we derive unified high probability upper bounds for all sub-Gaussian data and also provide matching minimax lower bounds in both Gaussian and binary cases. Due to the close connection of graphon to stochastic block models, an immediate consequence of our general results is a minimax rate-optimal estimator for sparse graphons.

Keywords. Biclustering; graphon; matrix completion; missing data; stochastic block models; sparse network.

1 Introduction

In a range of important data analytic scenarios, we encounter matrices with biclustering structures. For instance, in gene expression studies, one can organize the rows of a data matrix to correspond to individual cancer patients and the columns to transcripts. Then the patients are expected to form groups according to different cancer subtypes and the genes are also expected to exhibit clustering effect according to the different pathways they belong to. Therefore, after appropriate reordering of the rows and the columns, the data matrix is expected to have a biclustering structure contaminated by noises [20]. Here, the observed gene expression levels are real numbers. In a different context, such a biclustering structure can also be present in network data. For example, stochastic block model (SBM for short) [13] is a popular model for exchangeable networks. In SBMs, the graph nodes are partitioned into k disjoint communities and the probability that any pair of nodes are connected is determined entirely by the community memberships of the nodes. Consequently, if one rearranges the

nodes from the same communities together in the graph adjacency matrix, then the mean adjacency matrix, where each off-diagonal entry equals the probability of an edge connecting the nodes represented by the corresponding row and column, also has a biclustering structure.

The goal of the present paper is to develop a unified theory for the estimation (and completion when there are missing entries) of matrices with biclustering structures. To this end, we propose to consider the following general model

$$X_{ij} = \theta_{ij} + \epsilon_{ij}, \quad i \in [n_1], j \in [n_2], \tag{1}$$

where for any positive integer m, we let $[m] = \{1, \ldots, m\}$. Here, for each (i, j), $\theta_{ij} = \mathbb{E}[X_{ij}]$ and ϵ_{ij} is an independent piece of mean zero sub-Gaussian noise. Moreover, we allow entries to be missing completely at random [29]. Thus, let E_{ij} be i.i.d. Bernoulli random variables with success probability $p \in (0, 1]$, and

$$\Omega = \{(i, j) : E_{ij} = 1\}. \tag{2}$$

Our final observations are

$$X_{ij}, \quad (i, j) \in \Omega. \tag{3}$$

To model the biclustering structure, we focus on the case where there are k_1 row clusters and k_2 column clusters, and the values of $\{\theta_{ij}\}$ are taken as constant if the rows and the columns belong to the same clusters. The goal is then to recover the signal matrix $\theta \in \mathbb{R}^{n_1 \times n_2}$ from the observations (3). To accomodate most interesting cases, especially the case of undirected networks, we shall also consider the case where the data matrix X is symmetric with zero diagonals. In such cases, we also require $X_{ij} = X_{ji}$ and $E_{ij} = E_{ji}$ for all $i \neq j$.

Main contributions In this paper, we propose a unified estimation procedure for partially observed data matrix generated from model (1) – (3). We establish high probability upper bounds for the mean squared errors of the resulting estimators. In addition, we show that these upper bounds are minimax rate-optimal in both the continuous case and the binary case by providing matching minimax lower bounds. Furthermore, SBM can be viewed as a special case of the symmetric version of (1). Thus, an immediate application of our results is the network completion problem for SBMs. With partially observed network edges, our method gives a rate-optimal estimator for the success matrix of the whole network in both the dense and the sparse regimes, which further leads to rate-optimal graphon estimation in both regimes.

Connection to the literature If only a low rank constraint is imposed on the mean matrix θ, then (1) – (3) becomes what is known in the literature as the matrix completion problem [28]. An impressive list of algorithms and theories have been developed for this problem, including but not limited to [5, 16, 6, 4, 3, 17, 27, 19]. In this paper, we investigate an alternative biclustering structural assumption for the matrix completion problem, which

was first proposed by John Hartigan [12]. Note that a biclustering structure automatically implies low-rankness. However, if one applies a low rank matrix completion algorithm directly in the current setting, the resulting estimator suffers an inferior error bound to the minimax rate-optimal one. Thus, a full exploitation of the biclustering structure is necessary, which is the focus of the current paper.

The results of our paper also imply rate-optimal estimation for sparse graphons. Previous results on graphon estimation include [1], [30], [26], [7] and the references therein. The minimax rates for dense graphon estimation were derived by [9]. During the time when this paper is written, we have become aware of an independent result on optimal sparse graphon estimation by [18].

Organization After a brief introduction to notation, the rest of the paper is organized as follows. In Section 2, we introduce the precise formulation of the problem and propose a constrained least squares estimator for the mean matrix θ. In Section 3, we show that the proposed estimator leads to minimax optimal performance for both Gaussian and binary data. Section 4 presents some extensions of our results to sparse graphon estimation and adaptation. The proofs of the main results are laid out in Section 5, with some auxiliary results deferred to the appendix.

Notation For a vector $z \in [k]^n$, define the set $z^{-1}(a) = \{i \in [n] : z(i) = a\}$ for $a \in [k]$. For a set S, $|S|$ denotes its cardinality and $\mathbf{1}_S$ denotes the indicator function. For a matrix $A = (A_{ij}) \in \mathbb{R}^{n_1 \times n_2}$, the ℓ_2 norm and ℓ_∞ norm are defined by $\|A\| = \sqrt{\sum_{ij} A_{ij}^2}$ and $\|A\|_\infty = \max_{ij} |A_{ij}|$, respectively. The inner product for two matrices A and B is $\langle A, B\rangle = \sum_{ij} A_{ij}B_{ij}$. Given a subset $\Omega \in [n_1] \times [n_2]$, we use the notation $\langle A, B\rangle_\Omega = \sum_{(i,j)\in\Omega} A_{ij}B_{ij}$ and $\|A\|_\Omega = \sqrt{\sum_{(i,j)\in\Omega} A_{ij}^2}$. Given two numbers $a, b \in \mathbb{R}$, we use $a \vee b = \max(a, b)$ and $a \wedge b = \min(a, b)$. The floor function $\lfloor a \rfloor$ is the largest integer no greater than a, and the ceiling function $\lceil a \rceil$ is the smallest integer no less than a. For two positive sequences $\{a_n\}, \{b_n\}$, $a_n \lesssim b_n$ means $a_n \le Cb_n$ for some constant $C > 0$ independent of n, and $a_n \asymp b_n$ means $a_n \lesssim b_n$ and $b_n \lesssim a_n$. The symbols $\mathbb{P}$ and $\mathbb{E}$ denote generic probability and expectation operators whose distribution is determined from the context.

2 Constrained least squares estimation

Recall the generative model defined in (1) and also the definition of the set Ω in (2) of the observed entries. As we have mentioned, throughout the paper, we assume that the ϵ_{ij}'s are independent sub-Gaussian noises with sub-Gaussianity parameter uniformly bounded from above by $\sigma > 0$. More precisely, we assume

$$\mathbb{E}e^{\lambda\epsilon_{ij}} \le e^{\lambda^2\sigma^2/2}, \quad \text{for all } i \in [n_1], j \in [n_2] \text{ and } \lambda \in \mathbb{R}. \tag{4}$$

We consider two types of biclustering structures. One is rectangular and asymmetric,

where we assume that the mean matrix belongs to the following parameter space

$$\Theta_{k_1k_2}(M) = \Big\{\theta = (\theta_{ij}) \in \mathbb{R}^{n_1\times n_2} : \theta_{ij} = Q_{z_1(i)z_2(j)}, z_1 \in [k_1]^{n_1}, z_2 \in [k_2]^{n_2}, \\ Q \in [-M, M]^{k_1\times k_2}\Big\}. \tag{5}$$

In other words, the mean values within each bicluster is homogenous, i.e., $\theta_{ij} = Q_{ab}$ if the ith row belongs to the ath row cluster and the jth column belong to the bth column cluster. The other type of structures we consider is the square and symmetric case. In this case, we impose symmetry requirement on the data generating process, i.e., $n_1 = n_2 = n$ and

$$X_{ij} = X_{ji},\ E_{ij} = E_{ji},\ \text{for all } i \neq j. \tag{6}$$

Since the case is mainly motivated by undirected network data where there is no edge linking any node to itself, we also assume $X_{ii} = 0$ for all $i \in [n]$. Finally, the mean matrix is assumed to belong to the following parameter space

$$\Theta^s_k(M) = \Big\{\theta = (\theta_{ij}) \in \mathbb{R}^{n\times n} : \theta_{ii} = 0,\ \theta_{ij} = \theta_{ji} = Q_{z(i)z(j)} \text{ for } i > j, z \in [k]^n, \\ Q = Q^T \in [-M, M]^{k\times k}\Big\}. \tag{7}$$

We proceed by assuming that we know the parameter space Θ which can be either $\Theta_{k_1k_2}(M)$ or $\Theta^s_k(M)$ and the rate p of an independent entry being observed, The issues of adaptation to unknown numbers of clusters and unknown observation rate p are addressed later in Section 4.2 and Section 4.3. Given Θ and p, we propose to estimate θ by the following program

$$\min_{\theta\in\Theta}\left\{\|\theta\|^2 - \frac{2}{p}\langle X, \theta\rangle_\Omega\right\}. \tag{8}$$

If we define

$$Y_{ij} = X_{ij}E_{ij}/p, \tag{9}$$

then (8) is equivalent to the following constrained least squares problem

$$\min_{\theta\in\Theta}\|Y - \theta\|^2, \tag{10}$$

and hence the name of our estimator. When the data is binary, $\Theta = \Theta^s_k(1)$ and $p = 1$, the problem specializes to estimating the mean adjacency matrix in stochastic block models, and the estimator defined as the solution to (10) reduces to the least squares estimator in [9].

3 Main results

In this section, we provide theoretical justifications of the constrained least squares estimator defined as the solution to (10). Our first result is the following universal high probability upper bounds.

Theorem 3.1. *For any global optimizer of (10) and any constant $C' > 0$, there exists a constant $C > 0$ only depending on C' such that*

$$\|\hat{\theta} - \theta\|^2 \leq C\frac{M^2 \vee \sigma^2}{p}\left(k_1k_2 + n_1 \log k_1 + n_2 \log k_2\right),$$

with probability at least $1-\exp\left(-C'\left(k_1k_2 + n_1 \log k_1 + n_2 \log k_2\right)\right)$ *uniformly over* $\theta \in \Theta_{k_1k_2}(M)$ *and all error distributions satisfying (4). For the symmetric parameter space* $\Theta_k^s(M)$*, the bound is simplified to*

$$\|\hat{\theta} - \theta\|^2 \leq C\frac{M^2 \vee \sigma^2}{p}\left(k^2 + n \log k\right),$$

with probability at least $1 - \exp\left(-C'\left(k^2 + n \log k\right)\right)$ *uniformly over* $\theta \in \Theta_k^s(M)$ *and all error distributions satisfying (4).*

When $(M^2 \vee \sigma^2)$ is bounded, the rate in Theorem 3.1 is $\left(k_1k_2 + n_1 \log k_1 + n_2 \log k_2\right)/p$ which can be decomposed into two parts. The part involving k_1k_2 reflects the number of parameters in the biclustering structure, while the part involving $(n_1 \log k_1 + n_2 \log k_2)$ results from the complexity of estimating the clustering structures of rows and columns. It is the price one needs to pay for not knowing the clustering information. In contrast, the minimax rate for matrix completion under low rank assumption would be $(n_1 \vee n_2)(k_1 \wedge k_2)/p$ [19, 24], since without any other constraint the biclustering assumption implies that the rank of the mean matrix θ is at most $k_1 \wedge k_2$. Therefore, we have $\left(k_1k_2 + n_1 \log k_1 + n_2 \log k_2\right)/p \ll (n_1 \vee n_2)(k_1 \wedge k_2)/p$ as long as both $n_1 \vee n_2$ and $k_1 \wedge k_2$ tend to infinity. Thus, by fully exploiting the biclustering structure, we obtain a better convergence rate than only using the low rank assumption.

In the rest of this section, we discuss two most representative cases, namely the Gaussian case and the symmetric Bernoulli case. The latter case is also known in the literature as stochastic block models.

The Gaussian case Specializing Theorem 3.1 to Gaussian random variables, we obtain the following result.

Corollary 3.1. *Assume* $\epsilon_{ij} \overset{iid}{\sim} N(0, \sigma^2)$ *and* $M \leq C_1\sigma$ *for some constant* $C_1 > 0$*. For any constant* $C' > 0$*, there exists some constant* C *only depending on* C_1 *and* C' *such that*

$$\|\hat{\theta} - \theta\|^2 \leq C\frac{\sigma^2}{p}\left(k_1k_2 + n_1 \log k_1 + n_2 \log k_2\right),$$

with probability at least $1-\exp\left(-C'\left(k_1k_2 + n_1 \log k_1 + n_2 \log k_2\right)\right)$ *uniformly over* $\theta \in \Theta_{k_1k_2}(M)$*. For the symmetric parameter space* $\Theta_k^s(M)$*, the bound is simplified to*

$$\|\hat{\theta} - \theta\|^2 \leq C\frac{\sigma^2}{p}\left(k^2 + n \log k\right),$$

with probability at least $1 - \exp\left(-C'\left(k^2 + n \log k\right)\right)$ *uniformly over* $\theta \in \Theta_k^s(M)$*.*

We now present a rate matching lower bound in the Gaussian model to show that the result of Corollary 3.1 is minimax optimal. To this end, we use $\mathbb{P}_{(\theta,\sigma^2,p)}$ to indicate the probability distribution of the model $X_{ij} \overset{ind}{\sim} N(\theta_{ij}, \sigma^2)$ with observation rate p.

Theorem 3.2. *Assume* $\frac{\sigma^2}{p}\left(\frac{k_1k_2}{n_1n_2} + \frac{\log k_1}{n_2} + \frac{\log k_2}{n_1}\right) \lesssim M^2$. *There exist some constants* $C, c > 0$, *such that*

$$\inf_{\hat{\theta}} \sup_{\theta \in \Theta_{k_1k_2}(M)} \mathbb{P}_{(\theta,\sigma^2,p)} \left(\|\hat{\theta} - \theta\|^2 > C\frac{\sigma^2}{p}\left(k_1k_2 + n_1 \log k_1 + n_2 \log k_2\right) \right) > c,$$

when $\log k_1 \asymp \log k_2$, *and*

$$\inf_{\hat{\theta}} \sup_{\theta \in \Theta_k^s(M)} \mathbb{P}_{(\theta,\sigma^2,p)} \left(\|\hat{\theta} - \theta\|^2 > C\frac{\sigma^2}{p}\left(k^2 + n \log k\right) \right) > c.$$

The symmetric Bernoulli case When the observed matrix is symmetric with zero diagonal and Bernoulli random variables as its super-diagonal entries, it can be viewed as the adjacency matrix of an undirected network and the problem of estimating its mean matrix with missing data can be viewed as a network completion problem. Given a partially observed Bernoulli adjacency matrix $\{X_{ij}\}_{(i,j)\in\Omega}$, one can predict the unobserved edges by estimating the whole mean matrix θ.

Given a symmetric adjacency matrix $X = X^T \in \{0,1\}^{n\times n}$ with zero diagonals, the stochastic block model [13] assumes $\{X_{ij}\}_{i>j}$ are independent Bernoulli random variables with mean $\theta_{ij} = Q_{z(i)z(j)} \in [0,1]$ with some matrix $Q \in [0,1]^{k\times k}$ and some label vector $z \in [k]^n$. In other words, the probability that there is an edge between the ith and the jth nodes only depends on their community labels $z(i)$ and $z(j)$. The following class then includes all possible mean matrices of stochastic block models with n nodes and k clusters and with edge probabilities uniformly bounded by ρ:

$$\Theta_k^+(\rho) = \left\{\theta \in [0,1]^{n\times n} : \theta_{ii} = 0, \theta_{ij} = \theta_{ji} = Q_{z(i)z(j)}, Q = Q^T \in [0,\rho]^{k\times k}, z \in [k]^n\right\}. \quad (11)$$

By the definition in (7), $\Theta_k^+(\rho) \subset \Theta_k^s(\rho)$.

Assume that we observe each edge independently with probability p, Theorem 3.1 leads to the following result.

Corollary 3.2. *Consider the optimization problem (10) with* $\Theta = \Theta_k^s(\rho)$. *For any global optimizer* $\hat{\theta}$ *and any constant* $C' > 0$, *there exists a constant* $C > 0$ *only depending on* C' *such that*

$$\|\hat{\theta} - \theta\|^2 \le C\frac{\rho}{p}\left(k^2 + n\log k\right),$$

with probability at least $1 - \exp\left(-C'\left(k^2 + n\log k\right)\right)$ *uniformly over* $\theta \in \Theta_k^s(\rho) \supset \Theta_k^+(\rho)$.

When $\rho = p = 1$, Corollary 3.2 implies Theorem 2.1 in [9]. A rate matching lower bound is given by the following theorem. We denote the probability distribution of a stochastic block model with mean matrix $\theta \in \Theta_k^+(\rho)$ and observation rate p by $\mathbb{P}_{(\theta,p)}$.

Theorem 3.3. *For stochastic block models, we have*

$$\inf_{\hat{\theta}} \sup_{\theta \in \Theta_k^+(\rho)} \mathbb{P}_{(\theta,p)} \left(\|\hat{\theta} - \theta\|^2 > C \left(\frac{\rho\left(k^2 + n\log k\right)}{p} \wedge \rho^2 n^2 \right) \right) > c,$$

for some constants $C, c > 0$.

The lower bound is the minimum of two terms. When $\rho \geq \frac{k^2+n\log k}{pn^2}$, the rate becomes $\frac{\rho(k^2+n\log k)}{p} \wedge \rho^2 n^2 \asymp \frac{\rho(k^2+n\log k)}{p}$. It is achieved by the constrained least squares estimator according to Corollary 3.2. When $\rho < \frac{k^2+n\log k}{pn^2}$, the rate is dominated by $\rho^2 n^2$. In this case, a trivial zero estimator achieves the minimax rate.

In the case of $p = 1$, a comparable result has been found independently by [18]. However, our result here is more general as it accommodates missing observations. Moreover, the general upper bounds in Theorem 3.1 even hold for networks with weighted edges.

4 Extensions

In this section, we extends the estimation procedure and the theory in Sections 2 and 3 toward three directions: sparse graphon estimation, adaptation to unknown numbers of row and column clusters, and adaptation to unknown observation rate.

4.1 Sparse graphon estimation

Consider a random graph with adjacency matrix $\{X_{ij}\} \in \{0,1\}^{n\times n}$, whose sampling procedure is determined by

$$(\xi_1, ..., \xi_n) \sim \mathbb{P}_\xi, \quad X_{ij}|(\xi_i, \xi_j) \sim \text{Bernoulli}(\theta_{ij}), \quad \text{where } \theta_{ij} = f(\xi_i, \xi_j). \tag{12}$$

For $i \in [n]$, $X_{ii} = \theta_{ii} = 0$. Conditioning on $(\xi_1, ..., \xi_n)$, $X_{ij} = X_{ji}$ is independent across $i > j$. The function f on $[0,1]^2$, which is assumed to be symmetric, is called a graphon. The concept of graphon is originated from graph limit theory [14, 22, 8, 21] and the studies of exchangeable arrays [2, 15]. It is the underlying nonparametric object that generates the random graph. Statistical estimation of graphon has been considered by [30, 26, 9, 10, 23] for dense networks. Using Corollary 3.2, we present a result for sparse graphon estimation.

Let us start with specifying the function class of graphons. Define the derivative operator by

$$\nabla_{jk} f(x,y) = \frac{\partial^{j+k}}{(\partial x)^j (\partial y)^k} f(x,y),$$

and we adopt the convention $\nabla_{00} f(x,y) = f(x,y)$. The Hölder norm is defined as

$$||f||_{\mathcal{H}_\alpha} = \max_{j+k\leq \lfloor\alpha\rfloor} \sup_{x,y\in\mathcal{D}} |\nabla_{jk} f(x,y)| + \max_{j+k=\lfloor\alpha\rfloor} \sup_{(x,y)\neq(x',y')\in\mathcal{D}} \frac{|\nabla_{jk} f(x,y) - \nabla_{jk} f(x',y')|}{||(x-x', y-y')||^{\alpha-\lfloor\alpha\rfloor}},$$

where $\mathcal{D} = \{(x, y) \in [0, 1]^2 : x \geq y\}$. Then, the sparse graphon class with Hölder smoothness α is defined by

$$\mathcal{F}_\alpha(\rho, L) = \{0 \leq f \leq \rho : \|f\|_{\mathcal{H}_\alpha} \leq L\sqrt{\rho}, f(x, y) = f(y, x) \text{ for all } x \in \mathcal{D}\},$$

where $L > 0$ is the radius of the class, which is assumed to be a constant. As argued in [9], it is sufficient to approximate a graphon with Hölder smoothness by a piecewise constant function. In the random graph setting, a piecewise constant function is the stochastic block model. Therefore, we can use the estimator defined by (10). Using Corollary 3.2, a direct bias-variance tradeoff argument leads to the following result. An independent finding of the same result is also made by [18].

Corollary 4.1. *Consider the optimization problem (10) where $Y_{ij} = X_{ij}$ and $\Theta = \Theta_k^s(M)$ with $k = \lceil n^{\frac{1}{1+\alpha\wedge 1}} \rceil$ and $M = \rho$. Given any global optimizer $\hat{\theta}$ of (10), we estimate f by $\hat{f}(\xi_i, \xi_j) = \hat{\theta}_{ij}$. Then, for any constant $C' > 0$, there exists a constant $C > 0$ only depending on C' and L such that*

$$\frac{1}{n^2} \sum_{i,j\in[n]} \left(\hat{f}(\xi_i, \xi_j) - f(\xi_i, \xi_j)\right)^2 \leq \rho \left(n^{-\frac{2\alpha}{\alpha+1}} + \frac{\log n}{n}\right),$$

with probability at least $1 - \exp(-C'(n^{\frac{1}{\alpha+1}} + n \log n))$ uniformly over $f \in \mathcal{F}_\alpha(\rho, L)$ and $\mathbb{P}_\xi$.

Corollary 4.1 implies an interesting phase transition phenomenon. When $\alpha \in (0, 1)$, the rate becomes $\rho(n^{-\frac{2\alpha}{\alpha+1}} + \frac{\log n}{n}) \asymp \rho n^{-\frac{2\alpha}{\alpha+1}}$, which is the typical nonparametric rate times a sparsity index of the network. When $\alpha \geq 1$, the rate becomes $\rho(n^{-\frac{2\alpha}{\alpha+1}} + \frac{\log n}{n}) \asymp \frac{\rho \log n}{n}$, which does not depend on the smoothness α. Corollary 4.1 extends Theorem 2.3 of [9] to the case $\rho < 1$. In [30], the graphon f is defined in a different way. Namely, they considered the setting where $(\xi_1, ..., \xi_n)$ are i.i.d. Unif$[0, 1]$ random variables under $\mathbb{P}_\xi$. Then, the adjacency matrix is generated with Bernoulli random variables having means $\theta_{ij} = \rho f(\xi_i, \xi_j)$ for a nonparametric graphon f satisfying $\int_0^1 \int_0^1 f(x, y)dxdy = 1$. For this setting, with appropriate smoothness assumption, we can estimate f by $\hat{f}(\xi_i, \xi_j) = \hat{\theta}_{ij}/\rho$. The rate of convergence would be $\rho^{-1}(n^{-\frac{2\alpha}{\alpha+1}} + \frac{\log n}{n})$.

4.2 Adaptation to unknown numbers of clusters

We now provide an adaptive procedure for estimating θ without assuming the knowledge of the numbers of row and column clusters. We give details on the procedure for the asymmetric parameter spaces $\Theta_{k_1 k_2}(M)$, and that for the symmetric parameter spaces $\Theta_k^s(M)$ can be obtained similarly.

To adapt to k_1 and k_2, we split the data into two halves. Namely, sample i.i.d. T_{ij} from Bernoulli($\frac{1}{2}$). Define $\Delta = \{(i, j) \in [n_1] \times n_2 : T_{ij} = 1\}$. Define $Y_{ij}^\Delta = 2X_{ij}E_{ij}T_{ij}/p$ and $Y_{ij}^{\Delta^c} = 2X_{ij}E_{ij}(1 - T_{ij})/p$ for all $(i, j) \in [n_1] \times [n_2]$. Then, for some given (k_1, k_2), the least squares estimators using Y^Δ and Y^{Δ^c} are given by

$$\hat{\theta}_{k_1 k_2}^\Delta = \operatorname*{argmin}_{\theta\in\Theta_{k_1 k_2}(M)} \|Y^\Delta - \theta\|^2, \quad \hat{\theta}_{k_1 k_2}^{\Delta^c} = \operatorname*{argmin}_{\theta\in\Theta_{k_1 k_2}(M)} \|Y^{\Delta^c} - \theta\|^2.$$

Select the number of clusters by

$$(\hat{k}_1, \hat{k}_2) = \underset{(k_1,k_2)\in[n_1]\times[n_2]}{\operatorname{argmin}} \|\hat{\theta}^{\Delta}_{k_1k_2} - Y^{\Delta^c}\|^2_{\Delta^c},$$

and define $\hat{\theta}^{\Delta} = \hat{\theta}^{\Delta}_{\hat{k}_1\hat{k}_2}$. Similarly, we can also define $\hat{\theta}^{\Delta^c}$ by validate the number of clusters using Y^{Δ}. The final estimator is given by

$$\hat{\theta}_{ij} = \begin{cases} \hat{\theta}^{\Delta^c}_{ij}, & (i,j) \in \Delta; \\ \hat{\theta}^{\Delta}_{ij}, & (i,j) \in \Delta^c. \end{cases}$$

Theorem 4.1. *For any constant $C' > 0$, there exists a constant $C > 0$ only depending on C' such that*

$$\|\hat{\theta} - \theta\|^2 \le C\frac{M^2 \vee \sigma^2}{p}\left(k_1k_2 + n_1 \log k_1 + n_2 \log k_2 + \frac{\log(n_1+n_2)}{p}\right),$$

with probability at least $1 - \exp\left(-C'\left(k_1k_2 + n_1 \log k_1 + n_2 \log k_2\right)\right) - n^{-C'}$ uniformly over $\theta \in \Theta_{k_1k_2}(M)$ and all error distributions satisfying (4).

Compared with Theorem 3.1, the rate given by Theorem 4.1 has an extra $p^{-1}\log(n_1+n_2)$ term. A sufficient condition for this extra term to be inconsequential is $p \gtrsim \frac{\log(n_1+n_2)}{n_1 \wedge n_2}$.

4.3 Adaptation to unknown observation rate

The estimator (10) depends on the knowledge of the observation rate p. When p is not too small, such a knowledge is not necessary for achieving the desired rates. Define

$$\hat{p} = \frac{\sum_{i=1}^{n_1}\sum_{j=1}^{n_2} E_{ij}}{n_1n_2} \tag{13}$$

for the asymmetric and

$$\hat{p} = \frac{\sum_{1\le i<j\le n} E_{ij}}{\frac{1}{2}n(n-1)} \tag{14}$$

for the symmetric case, and redefine

$$Y_{ij} = X_{ij}E_{ij}/\hat{p} \tag{15}$$

where the actual definition of $\hat{p}$ is chosen between (13) and (14) depending on whether one is dealing with the asymmetric or symmetric parameter space. Then we have the following result for the solution to (10) with Y redefined by (15).

Theorem 4.2. *For $\Theta = \Theta_{k_1k_2}(M)$, suppose for some absolute constant $C_1 > 0$,*

$$p \ge C_1\frac{[\log(n_1+n_2)]^2}{k_1k_2 + n_1\log k_1 + n_2 \log k_2}.$$

Let $\hat{\theta}$ be the solution to (10) with Y defined as in (15). Then for any constant $C' > 0$, there exists a constant $C > 0$ only depending on C' and C_1 such that

$$\|\hat{\theta} - \theta\|^2 \leq C\frac{M^2 \vee \sigma^2}{p}\left(k_1k_2 + n_1 \log k_1 + n_2 \log k_2\right),$$

with probability at least $1 - (n_1n_2)^{-C'}$ uniformly over $\theta \in \Theta$ and all error distributions satisfying (4).

For $\Theta = \Theta_k^s(M)$, the same result holds if we replace n_1 and n_2 with n and k_1 and k_2 with k in the foregoing statement.

5 Proofs

5.1 Proof of Theorem 3.1

Below, we focus on the proof for the asymmetric parameter space $\Theta_{k_1k_2}(M)$. The result for the symmetric parameter space $\Theta_k^s(M)$ can be obtained by letting $k_1 = k_2$ and by taking care of the diagonal entries. Since $\hat{\theta} \in \Theta_{k_1k_2}(M)$, there exists $\hat{z}_1 \in [k_1]^{n_1}$, $\hat{z}_2 \in [k_2]^{n_2}$ and $\hat{Q} \in [-M, M]^{k_1 \times k_2}$ such that $\hat{\theta}_{ij} = \hat{Q}_{\hat{z}_1(i)\hat{z}_2(j)}$. For this $(\hat{z}_1, \hat{z}_2)$, we define a matrix $\tilde{\theta}$ by

$$\tilde{\theta}_{ij} = \frac{1}{|\hat{z}_1^{-1}(a)||\hat{z}_2^{-1}(b)|}\sum_{(i,j)\in\hat{z}_1^{-1}(a)\times\hat{z}_2^{-1}(b)}\theta_{ij},$$

for any $(i, j) \in \hat{z}_1^{-1}(a) \times \hat{z}_2^{-1}(b)$ and any $(a, b) \in [k_1] \times [k_2]$. To facilitate the proof, we need to following three lemmas, whose proofs are given in the supplementary material.

Lemma 5.1. *For any constant $C' > 0$, there exists a constant $C_1 > 0$ only depending on C', such that*

$$\|\hat{\theta} - \tilde{\theta}\|^2 \leq C_1\frac{M^2 \vee \sigma^2}{p}(k_1k_2 + n_1 \log k_1 + n_2 \log k_2),$$

with probability at least $1 - \exp(-C'(n_1 \log k_1 + n_2 \log k_2))$.

Lemma 5.2. *For any constant $C' > 0$, there exists a constant $C_2 > 0$ only depending on C', such that the inequality $\|\tilde{\theta} - \theta\|^2 \geq C_2(M^2 \vee \sigma^2)(n_1 \log k_1 + n_2 \log k_2)/p$ implies*

$$\left|\left\langle \frac{\tilde{\theta} - \theta}{\|\tilde{\theta} - \theta\|}, Y - \theta \right\rangle\right| \leq \sqrt{C_2\frac{M^2 \vee \sigma^2}{p}(k_1k_2 + n_1 \log k_1 + n_2 \log k_2)},$$

with probability at least $1 - \exp(-C'(k_1k_2 + n_1 \log k_1 + n_2 \log k_2))$.

Lemma 5.3. *For any constant $C' > 0$, there exists a constant $C_3 > 0$ only depending on C', such that*

$$\left|\left\langle \hat{\theta} - \tilde{\theta}, Y - \theta \right\rangle\right| \leq C_3\frac{M^2 \vee \sigma^2}{p}(k_1k_2 + n_1 \log k_1 + n_2 \log k_2),$$

with probability at least $1 - \exp(-C'(k_1k_2 + n_1 \log k_1 + n_2 \log k_2))$.

Proof of Theorem 3.1. Applying union bound, the results of Lemma 5.1-5.3 hold with probability at least $1 - 3\exp\left(-C'(k_1k_2 + n_1 \log k_1 + n_2 \log k_2)\right)$. We consider the following two cases.

Case 1: $\|\tilde{\theta}-\theta\|^2 \le C_2(M^2 \vee \sigma^2)(k_1k_2 + n_1 \log k_1 + n_2 \log k_2)/p$.
Then we have

$$\|\hat{\theta}-\theta\|^2 \le 2\|\hat{\theta}-\tilde{\theta}\|^2 + 2\|\tilde{\theta}-\theta\|^2 \le 2(C_1+C_2)\frac{M^2 \vee \sigma^2}{p}(k_1k_2 + n_1 \log k_1 + n_2 \log k_2)$$

by Lemma 5.1.

Case 2: $\|\tilde{\theta}-\theta\|^2 > C_2(M^2 \vee \sigma^2)(k_1k_2 + n_1 \log k_1 + n_2 \log k_2)/p$.
By the definition of the estimator, we have $\|\hat{\theta}-Y\|^2 \le \|\theta - Y\|^2$. After rearrangement, we have

$$\begin{aligned}
\|\hat{\theta}-\theta\|^2 &\le 2\left\langle \hat{\theta}-\theta, Y-\theta \right\rangle \\
&= 2\left\langle \hat{\theta}-\tilde{\theta}, Y-\theta \right\rangle + 2\left\langle \tilde{\theta}-\theta, Y-\theta \right\rangle \\
&\le 2\left\langle \hat{\theta}-\tilde{\theta}, Y-\theta \right\rangle + 2\|\tilde{\theta}-\theta\| \left| \left\langle \frac{\tilde{\theta}-\theta}{\|\tilde{\theta}-\theta\|}, Y-\theta \right\rangle \right| \\
&\le 2\left\langle \hat{\theta}-\tilde{\theta}, Y-\theta \right\rangle + 2(\|\tilde{\theta}-\hat{\theta}\| + \|\hat{\theta}-\theta\|) \left| \left\langle \frac{\tilde{\theta}-\theta}{\|\tilde{\theta}-\theta\|}, Y-\theta \right\rangle \right| \\
&\le 2(C_2 + C_3 + \sqrt{C_1C_2})\frac{M^2 \vee \sigma^2}{p}(k_1k_2 + n_1 \log k_1 + n_2 \log k_2) + \frac{1}{2}\|\hat{\theta}-\theta\|^2,
\end{aligned}$$

which leads to the bound

$$\|\hat{\theta}-\theta\|^2 \le 4(C_2 + C_3 + \sqrt{C_1C_2})\frac{M^2 \vee \sigma^2}{p}(k_1k_2 + n_1 \log k_1 + n_2 \log k_2).$$

Combining the two cases, we have

$$\|\hat{\theta}-\theta\|^2 \le C\frac{M^2 \vee \sigma^2}{p}(k_1k_2 + n_1 \log k_1 + n_2 \log k_2),$$

with probability at least $1 - 3\exp\left(-C'(k_1k_2 + n_1 \log k_1 + n_2 \log k_2)\right)$ for $C = 4(C_2 + C_3 + \sqrt{C_1C_2}) \vee 2(C_1 + C_2)$. □

5.2 Proof of Theorem 4.1

We first present a lemma for the tail behavior of sum of independent products of sub-Gaussian and Bernoulli random variables. Its proof is given in the supplementary material.

Lemma 5.4. *Let $\{X_i\}$ be independent sub-Gaussian random variables with mean $\theta_i \in [-M, M]$ and $\mathbb{E}e^{\lambda(X_i-\theta_i)} \le e^{\lambda^2\sigma^2/2}$. Let $\{E_i\}$ be independent Bernoulli random variables with mean p. Assume $\{X_i\}$ and $\{E_i\}$ are all independent. Then for $|\lambda| \le p/(M \vee \sigma)$ and $Y_i = X_iE_i/p$, we have*

$$\mathbb{E}e^{\lambda(Y_i-\theta_i)} \le 2e^{(M^2+2\sigma^2)\lambda^2/p}.$$

Moreover, for $\sum_{i=1}^n c_i^2 = 1$,

$$\mathbb{P}\left\{\left|\sum_{i=1}^n c_i(Y_i-\theta_i)\right| \geq t\right\} \leq 4\exp\left\{-\min\left(\frac{pt^2}{4(M^2+2\sigma^2)}, \frac{pt}{2(M\vee\sigma)\|c\|_\infty}\right)\right\} \tag{16}$$

for any $t>0$.

Proof of Theorem 4.1. By the definition of $(\hat{k}_1,\hat{k}_2)$, we have $\|\hat{\theta}^{\Delta}_{\hat{k}_1\hat{k}_2} - Y^{\Delta^c}\|^2_{\Delta^c} \leq \|\hat{\theta}^{\Delta}_{k_1k_2} - Y^{\Delta^c}\|^2_{\Delta^c}$. After rearrangement, we have

$$\begin{aligned}
\|\hat{\theta}^{\Delta}_{\hat{k}_1\hat{k}_2} - \theta\|^2_{\Delta^c} &\leq \|\hat{\theta}^{\Delta}_{k_1k_2} - \theta\|^2_{\Delta^c} + 2\|\hat{\theta}^{\Delta}_{\hat{k}_1\hat{k}_2} - \hat{\theta}^{\Delta}_{k_1k_2}\|_{\Delta^c}\left\langle \frac{\hat{\theta}^{\Delta}_{\hat{k}_1\hat{k}_2} - \hat{\theta}^{\Delta}_{k_1k_2}}{\|\hat{\theta}^{\Delta}_{\hat{k}_1\hat{k}_2} - \hat{\theta}^{\Delta}_{k_1k_2}\|_{\Delta^c}}, Y^{\Delta^c} - \theta\right\rangle_{\Delta^c} \\
&\leq \|\hat{\theta}^{\Delta}_{k_1k_2} - \theta\|^2_{\Delta^c} + 2\|\hat{\theta}^{\Delta}_{\hat{k}_1\hat{k}_2} - \hat{\theta}^{\Delta}_{k_1k_2}\|_{\Delta^c}\max_{(l_1,l_2)\in[n_1]\times[n_2]}\left|\left\langle \frac{\hat{\theta}^{\Delta}_{l_1l_2} - \hat{\theta}^{\Delta}_{k_1k_2}}{\|\hat{\theta}^{\Delta}_{l_1l_2} - \hat{\theta}^{\Delta}_{k_1k_2}\|_{\Delta^c}}, Y^{\Delta^c} - \theta\right\rangle_{\Delta^c}\right|.
\end{aligned}$$

By Lemma 5.4 and the independence structure, we have

$$\max_{(l_1,l_2)\in[n_1]\times[n_2]}\left|\left\langle \frac{\hat{\theta}^{\Delta}_{l_1l_2} - \hat{\theta}^{\Delta}_{k_1k_2}}{\|\hat{\theta}^{\Delta}_{l_1l_2} - \hat{\theta}^{\Delta}_{k_1k_2}\|_{\Delta^c}}, Y^{\Delta^c} - \theta\right\rangle_{\Delta^c}\right| \leq C(M\vee\sigma)\frac{\log(n_1+n_2)}{p},$$

with probability at least $1-n^{-C'}$. Using triangle inequality and Cauchy-Schwarz inequality, we have

$$\|\hat{\theta}^{\Delta}_{\hat{k}_1\hat{k}_2} - \theta\|^2_{\Delta^c} \leq \frac{3}{2}\|\hat{\theta}^{\Delta}_{k_1k_2} - \theta\|^2_{\Delta^c} + \frac{1}{2}\|\hat{\theta}^{\Delta}_{\hat{k}_1\hat{k}_2} - \theta\|^2_{\Delta^c} + 4C^2(M^2\vee\sigma^2)\left(\frac{\log(n_1+n_2)}{p}\right)^2.$$

By rearranging the above inequality, we have

$$\|\hat{\theta}^{\Delta}_{\hat{k}_1\hat{k}_2} - \theta\|^2_{\Delta^c} \leq 3\|\hat{\theta}^{\Delta}_{k_1k_2} - \theta\|^2_{\Delta^c} + 8C^2(M^2\vee\sigma^2)\left(\frac{\log(n_1+n_2)}{p}\right)^2.$$

A symmetric argument leads to

$$\|\hat{\theta}^{\Delta^c}_{\hat{k}_1\hat{k}_2} - \theta\|^2_{\Delta} \leq 3\|\hat{\theta}^{\Delta^c}_{k_1k_2} - \theta\|^2_{\Delta} + 8C^2(M^2\vee\sigma^2)\left(\frac{\log(n_1+n_2)}{p}\right)^2.$$

Summing up the above two inequalities, we have

$$\|\hat{\theta} - \theta\|^2 \leq 3\|\hat{\theta}^{\Delta}_{k_1k_2} - \theta\|^2 + 3\|\hat{\theta}^{\Delta^c}_{k_1k_2} - \theta\|^2 + 16C^2(M^2\vee\sigma^2)\left(\frac{\log(n_1+n_2)}{p}\right)^2.$$

Using Theorem 3.1 to bound $\|\hat{\theta}^{\Delta}_{k_1k_2} - \theta\|^2$ and $\|\hat{\theta}^{\Delta^c}_{k_1k_2} - \theta\|^2$, the proof is complete. □

5.3 Proof of Theorem 4.2

Recall the augmented data $Y_{ij} = X_{ij}E_{ij}/p$. Define $\mathcal{Y}_{ij} = X_{ij}E_{ij}/\hat{p}$. Let us give two lemmas to facilitate the proof.

Lemma 5.5. *Assume $p \gtrsim \frac{\log(n_1+n_2)}{n_1 n_2}$. For any $C' > 0$, there is some constant $C > 0$ such that*

$$\|Y - \mathcal{Y}\|^2 \le C\left[M^2 + \sigma^2 \log(n_1+n_2)\right] \frac{\log(n_1+n_2)}{p^2},$$

with probability at least $1 - (n_1 n_2)^{-C'}$.

Lemma 5.6. *The inequalities in Lemma 5.1-5.3 continue to hold with bounds*

$$C_1 \frac{M^2 \vee \sigma^2}{p}(k_1 k_2 + n_1 \log k_1 + n_2 \log k_2) + 2\|Y - \mathcal{Y}\|^2,$$

$$\sqrt{C_2 \frac{M^2 \vee \sigma^2}{p}(k_1 k_2 + n_1 \log k_1 + n_2 \log k_2)} + \|Y - \mathcal{Y}\|,$$

and

$$C_3 \frac{M^2 \vee \sigma^2}{p}(k_1 k_2 + n_1 \log k_1 + n_2 \log k_2) + \|\hat{\theta} - \tilde{\theta}\|\|Y - \mathcal{Y}\|,$$

respectively.

Proof of Theorem 4.2. The proof is similar to that of Theorem 3.1. We only need to replace Lemma 5.1-5.3 by Lemma 5.5 and Lemma 5.6 to get the desired result. □

5.4 Proofs of Theorem 3.2 and Theorem 3.3

This section gives proofs of the minimax lower bounds. We first introduce some notation. For any probability measures $\mathbb{P}, \mathbb{Q}$, define the Kullback–Leibler divergence by $D(\mathbb{P}||\mathbb{Q}) = \int \left(\log \frac{d\mathbb{P}}{d\mathbb{Q}}\right) d\mathbb{P}$. The chi-squared divergence is defined by $\chi^2(\mathbb{P}||\mathbb{Q}) = \int \left(\frac{d\mathbb{P}}{d\mathbb{Q}}\right) d\mathbb{P} - 1$. The main tool we will use is the following proposition.

Proposition 5.1. *Let (Ξ, ℓ) be a metric space and $\{\mathbb{P}_\xi : \xi \in \Xi\}$ be a collection of probability measures. For any totally bounded $T \subset \Xi$, define the Kullback-Leibler diameter and the chi-squared diameter of T by*

$$d_{\mathrm{KL}}(T) = \sup_{\xi,\xi' \in T} D(\mathbb{P}_\xi||\mathbb{P}_{\xi'}), \quad d_{\chi^2}(T) = \sup_{\xi,\xi' \in T} \chi^2(\mathbb{P}_\xi||\mathbb{P}_{\xi'}).$$

Then

$$\inf_{\hat{\xi}} \sup_{\xi \in \Xi} \mathbb{P}_\xi \left\{ \ell^2\left(\hat{\xi}(X), \xi\right) \ge \frac{\epsilon^2}{4} \right\} \ge 1 - \frac{d_{\mathrm{KL}}(T) + \log 2}{\log \mathcal{M}(\epsilon, T, \ell)}, \tag{17}$$

$$\inf_{\hat{\xi}} \sup_{\xi \in \Xi} \mathbb{P}_\xi \left\{ \ell^2\left(\hat{\xi}(X), \xi\right) \ge \frac{\epsilon^2}{4} \right\} \ge 1 - \frac{1}{\mathcal{M}(\epsilon, T, \ell)} - \sqrt{\frac{d_{\chi^2}(T)}{\mathcal{M}(\epsilon, T, \ell)}}, \tag{18}$$

for any $\epsilon > 0$, where the packing number $\mathcal{M}(\epsilon, T, \ell)$ is the largest number of points in T that are at least ϵ away from each other.

The inequality (17) is the classical Fano's inequality. The version we present here is by [31]. The inequality (18) is a generalization of the classical Fano's inequality by using chi-squared divergence instead of KL divergence. It is due to [11].

The following proposition bounds the KL divergence and the chi-squared divergence for both Gaussian and Bernoulli models.

Proposition 5.2. *For the Gaussian model, we have*

$$D\left(\mathbb{P}_{(\theta,\sigma^2,p)}||\mathbb{P}_{(\theta',\sigma^2,p)}\right) \le \frac{p}{2\sigma^2}\|\theta-\theta'\|^2, \quad \chi^2\left(\mathbb{P}_{(\theta,\sigma^2,p)}||\mathbb{P}_{(\theta',\sigma^2,p)}\right) \le \exp\left(\frac{p}{\sigma^2}\|\theta-\theta'\|^2\right)-1.$$

For the Bernoulli model with any $\theta,\theta' \in [\rho/2, 3\rho/4]^{n_1\times n_2}$, we have

$$D\left(\mathbb{P}_{(\theta,p)}||\mathbb{P}_{(\theta',p)}\right) \le \frac{8p}{\rho}\|\theta-\theta'\|^2, \quad \chi^2\left(\mathbb{P}_{(\theta,p)}||\mathbb{P}_{(\theta',p)}\right) \le \exp\left(\frac{8p}{\rho}\|\theta-\theta'\|^2\right)-1.$$

Finally, we need the following Varshamov–Gilbert bound. The version we present here is due to [25, Lemma 4.7].

Lemma 5.7. *There exists a subset $\{\omega_1,...,\omega_N\} \subset \{0,1\}^d$ such that*

$$H(\omega_i,\omega_j) \triangleq \|\omega_i-\omega_j\|^2 \ge \frac{d}{4}, \quad \text{for any } i\ne j\in[N], \tag{19}$$

for some $N \ge \exp(d/8)$.

Proof of Theorem 3.2. We focus on the proof for the asymmetric parameter space $\Theta_{k_1k_2}(M)$. The result for the symmetric parameter space $\Theta_k^s(M)$ can be obtained by letting $k_1=k_2$ and by taking care of the diagonal entries. Let us assume n_1/k_1 and n_2/k_2 are integers without loss of generality. We first derive the lower bound for the nonparametric rate $\sigma^2k_1k_2/p$. Let us fix the labels by $z_1(i)=\lceil ik_1/n_1\rceil$ and $z_2(j)=\lceil jk_2/n_2\rceil$. For any $\omega\in\{0,1\}^{k_1\times k_2}$, define

$$Q^\omega_{ab} = c\sqrt{\frac{\sigma^2k_1k_2}{pn_1n_2}}\omega_{ab}. \tag{20}$$

By Lemma 5.7, there exists some $T\subset\{0,1\}^{k_1k_2}$ such that $|T|\ge\exp(k_1k_2/8)$ and $H(\omega,\omega')\ge k_1k_2/4$ for any $\omega,\omega\in T$ and $\omega\ne\omega'$. We construct the subspace

$$\Theta(z_1,z_2,T) = \left\{\theta\in\mathbb{R}^{n_1\times n_2} : \theta_{ij} = Q^\omega_{z_1(i)z_2(j)}, \omega\in T\right\}.$$

By Proposition 5.2, we have

$$\sup_{\theta,\theta'\in\Theta(z_1,z_2,T)} \chi^2\left(\mathbb{P}_{(\theta,\sigma^2,p)}||\mathbb{P}_{(\theta',\sigma^2,p)}\right) \le \exp\left(c^2k_1k_2\right).$$

For any two different θ and θ' in $\Theta(z_1,z_2,T)$ associated with $\omega,\omega'\in T$, we have

$$\|\theta-\theta'\|^2 \ge \frac{c^2\sigma^2}{p}H(\omega,\omega') \ge \frac{c^2\sigma^2}{4p}k_1k_2.$$

Therefore, $\mathcal{M}\left(\sqrt{\frac{c^2\sigma^2}{4p}k_1k_2}, \Theta(z_1, z_2, T), \|\cdot\|\right) \geq \exp(k_1k_2/8)$. Using (18) with an appropriate c, we have obtained the rate $\frac{\sigma^2}{p}k_1k_2$ in the lower bound.

Now let us derive the clustering rate $\sigma^2 n_2 \log k_2/p$. Let us pick $\omega_1, ..., \omega_{k_2} \in \{0,1\}^{k_1}$ such that $H(\omega_a, \omega_b) \geq \frac{k_1}{4}$ for all $a \neq b$. By Lemma 5.7, this is possible when $\exp(k_1/8) \geq k_2$. Then, define

$$Q_{*a} = c\sqrt{\frac{\sigma^2 n_2 \log k_2}{pn_1n_2}}\omega_a. \tag{21}$$

Define z_1 by $z_1(i) = \lceil ik_1/n_1 \rceil$. Fix Q and z_1 and we are gong to let z_2 vary. Select a set $\mathcal{Z}_2 \subset [k_2]^{n_2}$ such that $|\mathcal{Z}_2| \geq \exp(Cn_2 \log k_2)$ and $H(z_2, z_2') \geq \frac{n_2}{6}$ for any $z_2, z_2' \in \mathcal{Z}_k$ and $z_2 \neq z_2'$. The existence of such $\mathcal{Z}_2$ is proved by [9]. Then, the subspace we consider is

$$\Theta(z_1, \mathcal{Z}_2, Q) = \left\{\theta \in \mathbb{R}^{n_1 \times n_2} : \theta_{ij} = Q_{z_1(i)z_2(j)}, z_2 \in \mathcal{Z}_2\right\}.$$

By Proposition 5.2, we have

$$\sup_{\theta,\theta' \in \Theta(z_1,\mathcal{Z}_2,Q)} D\left(\mathbb{P}_{(\theta,\sigma^2,p)} || \mathbb{P}_{(\theta',\sigma^2,p)}\right) \leq c^2 n_2 \log k_2.$$

For any two different θ and θ' in $\Theta(z_1, \mathcal{Z}_2, Q)$ associated with $z_2, z_2' \in \mathcal{Z}_2$, we have

$$\|\theta - \theta'\|^2 = \sum_{j=1}^{n_2} \|\theta_{*z_2(j)} - \theta'_{*z_2(j)}\|^2 \geq H(z_2, z_2')\frac{c^2\sigma^2 n_2 \log k_2}{pn_1n_2}\frac{n_1}{4} \geq \frac{c^2\sigma^2 n_2 \log k_2}{24p}.$$

Therefore, $\mathcal{M}\left(\sqrt{\frac{c^2\sigma^2 n_2 \log k_2}{24p}}, \Theta(z_1, \mathcal{Z}_2, Q), \|\cdot\|\right) \geq \exp(Cn_2 \log k_2)$. Using (17) with some appropriate c, we obtain the lower bound $\frac{\sigma^2 n_2 \log k_2}{p}$.

A symmetric argument gives the rate $\frac{\sigma^2 n_1 \log k_1}{p}$. Combining the three parts using the same argument in [9], the proof is complete. □

Proof of Theorem 3.3. The proof is similar to that of Theorem 3.2. The only differences are (20) replaced by

$$Q^{\omega}_{ab} = \frac{1}{2}\rho + \left(c\sqrt{\frac{\rho k^2}{pn^2}} \wedge \frac{1}{2}\rho\right)\omega_{ab}$$

and (21) replaced by

$$Q_{*a} = \frac{1}{2}\rho + \left(c\sqrt{\frac{\rho \log k}{pn}} \wedge \frac{1}{2}\rho\right)\omega_a.$$

It is easy to check that the constructed subspaces are subsets of $\Theta^+_k(\rho)$. Then, a symmetric modification of the proof of Theorem 3.2 leads to the desired conclusion. □

A Proofs of auxiliary results

In this section, we give proofs of Lemma 5.1-5.5. We first introduce some notation. Define the set

$$\mathcal{Z}_{k_1k_2} = \{z = (z_1, z_2) : z_1 \in [k_1]^{n_1}, z_2 \in [k_2]^{n_2}\}.$$

For an $n_1 \times n_2$ and some $z = (z_1, z_2) \in \mathcal{Z}_{k_1k_2}$, define

$$\bar{R}_{ab}(z) = \frac{1}{|z_1^{-1}(a)||z_2^{-1}(b)|} \sum_{(i,j)\in z_1^{-1}(a)\times z_2^{-1}(b)} R_{ij},$$

for all $a \in [k_1], b \in [k_2]$. To facilitate the proof, we need the following two results.

Proposition A.1. *For the estimator $\hat{\theta}_{ij} = \hat{Q}_{\hat{z}_1(i)\hat{z}_2(j)}$, we have*

$$\hat{Q}_{ab} = \operatorname{sign}(\bar{Y}_{ab}(\hat{z}))\left(|\bar{Y}_{ab}(\hat{z})| \wedge M\right),$$

for all $a \in [k_1], b \in [k_2]$.

Lemma A.1. *Under the setting of Lemma 5.4, define $S = \frac{1}{\sqrt{n}}\sum_{i=1}^n (Y_i - \theta_i)$ and $\tau = 2(M^2 + 2\sigma^2)/(M \vee \sigma)$. Then we have the following results:*

a. Let $T = S\mathbf{1}\{|S| \le \tau\sqrt{n}\}$, then $\mathbb{E}e^{pT^2/(8(M^2+2\sigma^2))} \le 5$;

b. Let $R = \tau\sqrt{n}|S|\mathbf{1}\{|S| > \tau\sqrt{n}\}$, then $\mathbb{E}e^{pR/(8(M^2+2\sigma^2))} \le 8$.

Proof. By (16),

$$\mathbb{P}\Big(|S| > t\Big) \le 4\exp\left\{-\min\left(\frac{pt^2}{4(M^2+2\sigma^2)}, \frac{\sqrt{n}pt}{2(M\vee\sigma)}\right)\right\}.$$

Then

$$\begin{aligned}
\mathbb{E}e^{\lambda T^2} &= \int_0^\infty \mathbb{P}\Big(e^{\lambda T^2} > u\Big)du \le 1 + \int_1^\infty \mathbb{P}\left(|T| > \sqrt{\frac{\log u}{\lambda}}\right)du \\
&= 1 + \int_1^{e^{\lambda\tau^2 n}} \mathbb{P}\left(|S| > \sqrt{\frac{\log u}{\lambda}}\right)du = 1 + 4\int_1^{e^{\lambda\tau^2 n}} u^{-p/(4\lambda(M^2+2\sigma^2))}du.
\end{aligned}$$

Choosing $\lambda = p/(8(M^2+2\sigma^2))$, we get $\mathbb{E}e^{pT^2/(8(M^2+2\sigma^2))} \le 5$. We proceed to prove the second claim.

$$\begin{aligned}
\mathbb{E}e^{\lambda R} &= \mathbb{P}(R=0) + \mathbb{P}(R>0)\mathbb{E}[e^{\lambda R}|R>0] \\
&= \mathbb{P}(R=0) + \mathbb{P}(R>0)\int_0^\infty \mathbb{P}(e^{\lambda R} > u|R>0)du \\
&= \mathbb{P}(R=0) + \int_0^\infty \mathbb{P}(e^{\lambda R} > u, R>0)du \\
&\le \mathbb{P}(R=0) + \mathbb{P}(R>0)e^{\lambda\tau^2 n} + \int_{e^{\lambda\tau^2 n}}^\infty \mathbb{P}(e^{\lambda R} > u)du \\
&\le 1 + 4e^{-p\tau^2 n/(4(M^2+2\sigma^2))+\lambda\tau^2 n} + \int_{e^{\lambda\tau^2 n}}^\infty \mathbb{P}(e^{\sqrt{n}\lambda\tau|S|} > u)du \\
&= 1 + 4e^{-p\tau^2 n/(4(M^2+3\sigma^2))+\lambda\tau^2 n} + 4\int_{e^{\lambda\tau^2 n}}^\infty u^{-p/(2\lambda\tau(M\vee\sigma))}du
\end{aligned}$$

Choosing $\lambda = p/(8(M^2+2\sigma^2))$, we get $\mathbb{E}e^{pR/(8(M^2+2\sigma^2))} \le 8$. □

Proof of Lemma 5.1. By the definitions of $\hat{\theta}_{ij}$ and $\tilde{\theta}_{ij}$ and Proposition A.1, we have

$$\hat{\theta}_{ij} - \tilde{\theta}_{ij} = \begin{cases} M - \bar{\theta}_{ab}(\hat{z}), & \text{if } \bar{Y}_{ab}(\hat{z}) \ge M; \\ \bar{Y}_{ab}(\hat{z}) - \bar{\theta}_{ab}(\hat{z}), & \text{if } -M \le \bar{Y}_{ab}(\hat{z}) < M; \\ -M - \bar{\theta}_{ab}(\hat{z}), & \text{if } \bar{Y}_{ab}(\hat{z}) < -M \end{cases}$$

for any $(i,j) \in \hat{z}_1^{-1}(a) \times \hat{z}_2^{-1}(b)$. Define $W = Y - \theta$, and it is easy to check that

$$|\hat{\theta}_{ij} - \tilde{\theta}_{ij}| \le |\bar{W}_{ab}(\hat{z})| \wedge 2M \le |\bar{W}_{ab}(\hat{z})| \wedge \tau,$$

where $\hat{z} = (\hat{z}_1, \hat{z}_2)$ and τ is defined in Lemma A.1. Then

$$\begin{aligned} \|\hat{\theta} - \tilde{\theta}\|^2 &\le \sum_{a\in[k_1], b\in[k_2]} \left|\hat{z}_1^{-1}(a)\right| \left|\hat{z}_2^{-1}(b)\right| \left(|\bar{W}_{ab}(\hat{z})| \wedge \tau\right)^2 \\ &\le \max_{z\in\mathcal{Z}_{k_1k_2}} \sum_{a\in[k_1], b\in[k_2]} \left|z_1^{-1}(a)\right| \left|z_2^{-1}(b)\right| \left(|\bar{W}_{ab}(z)| \wedge \tau\right)^2. \end{aligned} \tag{22}$$

For any $a \in [k_1], b \in [k_2]$ and $z_1 \in [k_1]^{n_1}$, $z_2 \in [k_2]^{n_2}$, define $n_1(a) = \left|z_1^{-1}(a)\right|$, $n_2(b) = \left|z_2^{-1}(b)\right|$ and $V_{ab}(z) = \sqrt{n_1(a)n_2(b)}|\bar{W}_{ab}(z)|\mathbf{1}\{|\bar{W}_{ab}(z)| \le \tau\}$. By Markov's inequality and Lemma A.1, we have

$$\begin{aligned} \mathbb{P}\left(\sum_{a\in[k_1], b\in[k_2]} V_{ab}^2(z) > t\right) &\le e^{-pt/(8(M^2+2\sigma^2))} \prod_{a\in[k_1], b\in[k_2]} e^{pV_{ab}^2(z)/(8(M^2+2\sigma^2))} \\ &\le \exp\left\{-\frac{pt}{8(M^2+2\sigma^2)} + k_1k_2\log 5\right\}. \end{aligned}$$

Applying union bound and using the fact that $\log|[k_1]^n| + \log|[k_2]^{n_2}| = n_1\log k_1 + n_2 \log k_2$,

$$\mathbb{P}\left(\max_{z\in\mathcal{Z}_{k_1k_2}} \sum_{a\in[k_1], b\in[k_2]} V_{ab}^2(z) > t\right) \le \exp\left\{-\frac{pt}{8(M^2+2\sigma^2)} + k_1k_2\log 5 + n_1\log k_1 + n_2\log k_2\right\}.$$

For any given constant $C' > 0$, we choose $t = C_1\frac{M^2\vee\sigma^2}{p}(k_1k_2 + n_1\log k_1 + n_2\log k_2)$ for some sufficiently large $C_1 > 0$ to obtain

$$\max_{z\in\mathcal{Z}_{k_1k_2}} \sum_{a\in[k_1], b\in[k_2]} V_{ab}^2(z) \le C_1\frac{M^2\vee\sigma^2}{p}(k_1k_2 + n_1\log k_1 + n_2\log k_2) \tag{23}$$

with probability at least $1 - \exp\left(-C'(k_1k_2 + n_1\log k_1 + n_2\log k_2)\right)$. Plugging (23) into (22), we complete the proof. □

Proof of Lemma 5.2. Note that

$$\tilde{\theta}_{ij} - \theta_{ij} = \sum_{a\in[k_1],b\in[k_2]} \bar{\theta}_{ab}(\hat{z})\mathbf{1}\{(i,j) \in \hat{z}_1^{-1}(a) \times \hat{z}_2^{-1}(b)\} - \theta_{ij}$$

is a function of $\hat{z}_1$ and $\hat{z}_2$. Then we have

$$\left|\sum_{ij} \frac{\tilde{\theta}_{ij} - \theta_{ij}}{\sqrt{\sum_{ij}(\tilde{\theta}_{ij} - \theta_{ij})^2}}(Y_{ij} - \theta_{ij})\right| \le \max_{z\in\mathcal{Z}_{k_1k_2}} \left|\sum_{ij} \gamma_{ij}(z)(Y_{ij} - \theta_{ij})\right|,$$

where

$$\gamma_{ij}(z) \propto \sum_{a\in[k_1],b\in[k_2]} \bar{\theta}_{ab}(z)\mathbf{1}\{(i,j) \in {z_1}^{-1}(a) \times {z_2}^{-1}(b)\} - \theta_{ij}$$

satisfies $\sum_{ij}\gamma_{ij}(z)^2 = 1$. Consider the event $\|\tilde{\theta} - \theta\|^2 \ge C_2(M^2 \vee \sigma^2)(k_1k_2 + n_1\log k_1 + n_2 \log k_2)/p$ for some C_2 to be specified later, we have

$$|\gamma_{ij}(z)| \le \frac{2M}{\|\tilde{\theta} - \theta\|} \le \sqrt{\frac{4M^2p}{C_2(M^2 \vee \sigma^2)(k_1k_2 + n_1\log k_1 + n_2\log k_2)}}.$$

By Lemma 5.4 and union bound, we have

$$\begin{aligned}
&\mathbb{P}\left(\max_{z\in\mathcal{Z}_{k_1k_2}} \left|\sum_{ij} \gamma_{ij}(z)(Y_{ij} - \theta_{ij})\right| > t\right) \\
\le\ & \sum_{z_1\in[k_1]^{n_1}, z_2\in[k_2]^{n_2}} \mathbb{P}\left(\left|\sum_{ij} \gamma_{ij}(z)(Y_{ij} - \theta_{ij})\right| > t\right) \\
\le\ & \exp\left(-C'\left(k_1k_2 + n_1\log k_1 + n_2\log k_2\right)\right),
\end{aligned}$$

by setting $t = \sqrt{C_2(M^2 \vee \sigma^2)(k_1k_2 + n_1\log k_1 + n_2\log k_2)/p}$ for some sufficiently large C_2 depending on C'. Thus, the lemma is proved. □

Proof of Lemma 5.3. By definition,

$$\begin{aligned}
&\left|\left\langle \hat{\theta} - \tilde{\theta}, Y - \theta\right\rangle\right| \\
=\ & \left|\sum_{a\in[k_1],b\in[k_2]} \Big(\mathrm{sign}(\bar{Y}_{ab}(\hat{z}))\left(|\bar{Y}_{ab}(\hat{z})| \wedge M\right) - \bar{\theta}_{ab}(\hat{z})\Big)\bar{W}_{ab}(\hat{z})|\hat{z}_1^{-1}(a)||\hat{z}_2^{-1}(b)|\right| \\
\le\ & \max_{z\in\mathcal{Z}_{k_1k_2}} \left|\sum_{a\in[k_1],b\in[k_2]} \Big(\mathrm{sign}(\bar{Y}_{ab}(z))\left(|\bar{Y}_{ab}(z)| \wedge M\right) - \bar{\theta}_{ab}(z)\Big)\bar{W}_{ab}(z)|z_1^{-1}(a)||z_2^{-1}(b)|\right|.
\end{aligned}$$

By definition, we have

$$\Big(\mathrm{sign}(\bar{Y}_{ab}(z))\left(|\bar{Y}_{ab}(z)| \wedge M\right) - \bar{\theta}_{ab}(z)\Big)\bar{W}_{ab}(z) \le |\bar{W}_{ab}(z)|^2 \wedge \tau|\bar{W}_{ab}(z)|.$$

For any fixed $z_1 \in [k_1]^{n_1}, z_2 \in [k_2]^{n_2}$, define $n_1(a) = |z_1^{-1}(a)|$ for $a \in [k_1]$, $n_2(b) = |z_1^{-1}(b)|$ for $b \in [k_2]$ and $V_{ab}(z) = \sqrt{n_1(a)n_2(b)}|\bar{W}_{ab}(z)|\mathbf{1}\{|\bar{W}_{ab}(z)| \le \tau\}$, $R_{ab}(z) = \tau n_1(a)n_2(b)|\bar{W}_{ab}(z)|\mathbf{1}\{|\bar{W}_{ab}(z)| > \tau\}$. Then

$$\left|\left\langle \hat{\theta} - \tilde{\theta}, Y - \theta \right\rangle\right| \le \max_{z \in \mathcal{Z}_{k_1 k_2}} \left\{ \sum_{a\in[k_1], b\in[k_2]} V_{ab}^2(z) + \sum_{a\in[k_1], b\in[k_2]} R_{ab}(z) \right\}.$$

Using Markov's inequality together with Lemma A.1, we have

$$\begin{aligned}
\mathbb{P}\left\{ \sum_{a\in[k_1], b\in[k_2]} V_{ab}^2(z) \ge t \right\} &\le e^{-pt/(8(M^2+2\sigma^2))} \prod_{a\in[k_1], b\in[k_2]} \mathbb{E} e^{pV_{ab}^2(z)/(8(M^2+2\sigma^2))} \\
&\le \exp\{-pt/(8(M^2+2\sigma^2)) + k_1 k_2 \log 5\}
\end{aligned}$$

$$\begin{aligned}
\mathbb{P}\left\{ \sum_{a\in[k_1], b\in[k_2]} R_{ab}(z) \ge t \right\} &\le e^{-pt/(8(M^2+2\sigma^2))} \prod_{a\in[k_1], b\in[k_2]} \mathbb{E} e^{pR_{ab}^2(z)/(8(M^2+2\sigma^2))} \\
&\le \exp\{-pt/(8(M^2+2\sigma^2)) + k_1 k_2 \log 8\}
\end{aligned}$$

Following the same argument in the proof of Lemma 5.1, a choice of $t = C_3(M^2 \vee \sigma^2)(k_1 k_2 + n_1 \log k_1 + n_2 \log k_2)/p$ for some sufficiently large $C_3 > 0$ will complete the proof. □

Proof of Lemma 5.4. When $|\lambda| \le p/(M \vee \sigma)$, $|\lambda\theta_i/p| \le 1$ and $\lambda^2\sigma^2/p^2 \le 1$. Then

$$\begin{aligned}
\mathbb{E} e^{\lambda(Y_i - \theta_i)} &= p\mathbb{E} e^{\lambda(X/p - \theta_i)} + (1-p)\mathbb{E} e^{-\lambda\theta_i} \\
&\le p e^{\frac{\lambda^2\sigma^2}{2p^2} + \frac{1-p}{p}\lambda\theta_i} + (1-p)e^{-\lambda\theta_i} \\
&\le p\left(1 + \frac{1-p}{p}\lambda\theta_i + \frac{2(1-p)^2}{p^2}\lambda^2\theta_i^2\right)\left(1 + \frac{\lambda^2\sigma^2}{p^2}\right) + (1-p)(1 - \lambda\theta_i + 2\lambda^2\theta_i^2) \\
&\le 1 + \frac{2(1-p)\theta_i^2 + \sigma^2}{p}\lambda^2 + \frac{1-p}{p^2}\lambda^3\theta_i\sigma^2 + \frac{2(1-p)^2}{p^3}\lambda^4\sigma^2\theta_i^2 + 2(1-p)\lambda^2\theta_i^2 \\
&\le 2 + (2M^2 + \sigma^2)\lambda^2/p + \lambda^3\theta_i\sigma^2/p^2 + 2\lambda^4\sigma^2\theta_i^2/p^3 \\
&\le 2 + (2M^2 + \sigma^2)\lambda^2/p + \lambda^2\sigma^2/p + 2\lambda^2\sigma^2/p \\
&\le 2 + (2M^2 + 4\sigma^2)\lambda^2/p \\
&\le 2e^{(M^2+2\sigma^2)\lambda^2/p}.
\end{aligned}$$

The second inequality is due to the fact that $e^x \le 1 + 2x$ for all $x \ge 0$ and $e^x \le 1 + x + 2x^2$ for all $|x| \le 1$. Then for $|\lambda|(M \vee \sigma)\|c\|_\infty \le p$, Markov inequality implies

$$\mathbb{P}\left(\sum_{i=1}^n c_i(Y_i - \theta_i) \ge t\right) \le 2\exp\left\{-\lambda t + \frac{\lambda^2}{p}(M^2 + 2\sigma^2)\right\}.$$

By choosing $\lambda = \min\left\{\frac{pt}{2(M^2+2\sigma^2)}, \frac{p}{(M\vee\sigma)\|c\|_\infty}\right\}$, we get (16). □

Proof of Lemma 5.5. By the definitions of Y and $\mathcal{Y}$, we have

$$\|Y - \mathcal{Y}\|^2 \le (\hat{p}^{-1} - p^{-1})^2 \max_{i,j} X_{ij}^2 \sum_{ij} E_{ij}.$$

Therefore, it is sufficient to bound the three terms. For the first term, we have

$$|\hat{p}^{-1} - p^{-1}| \le |\hat{p}^{-1} - p^{-1}| \frac{|\hat{p} - p|}{p} + \frac{|\hat{p} - p|}{p^2},$$

which leads to

$$|\hat{p}^{-1} - p^{-1}| \le \left(1 - \frac{|\hat{p} - p|}{p}\right)^{-1} \frac{|\hat{p} - p|}{p^2}. \tag{24}$$

Bernstein's inequality implies $|\hat{p} - p|^2 \le C\frac{p\log(n_1+n_2)}{n_1 n_2}$ with probability at least $1 - (n_1 n_2)^{-C'}$ under the assumption that $p \gtrsim \frac{\log(n_1+n_2)}{n_1 n_2}$. Plugging the bound into (24), we get

$$(\hat{p}^{-1} - p^{-1})^2 \le C_1 \frac{\log(n_1 + n_2)}{p^3 n_1 n_2}.$$

The second term can be bounded by a union bound with the sub-Gaussian tail assumption of each X_{ij}. That is,

$$\max_{i,j} X_{ij}^2 \le C_2(M^2 + \sigma^2 \log(n_1 + n_2)),$$

with probability at least $1 - (n_1 n_2)^{-C'}$. Finally, using Bernstein's inequality again, the third term is bounded as

$$\sum_{ij} E_{ij} \le C_3 n_1 n_2 \left(p + \sqrt{\frac{p\log(n_1 + n_2)}{n_1 n_2}}\right) \le C_3' n_1 n_2 p,$$

with probability at least $1 - (n_1 n_2)^{-C'}$ under the assumption that $p \gtrsim \frac{\log(n_1+n_2)}{n_1 n_2}$. Combining the three bounds, we have obtained the desired conclusion. □

Proof of Lemma 5.6. For the second and the third bounds, we use

$$\left|\left\langle \frac{\tilde{\theta} - \theta}{\|\tilde{\theta} - \theta\|}, \mathcal{Y} - \theta \right\rangle\right| \le \left|\left\langle \frac{\tilde{\theta} - \theta}{\|\tilde{\theta} - \theta\|}, Y - \theta \right\rangle\right| + \|\mathcal{Y} - Y\|,$$

and

$$\left|\left\langle \hat{\theta} - \tilde{\theta}, \mathcal{Y} - \theta \right\rangle\right| \le \left|\left\langle \hat{\theta} - \tilde{\theta}, Y - \theta \right\rangle\right| + \|\hat{\theta} - \tilde{\theta}\|\|Y - \mathcal{Y}\|,$$

followed by the original proofs of Lemma 5.2 and Lemma 5.3. To prove the first bound, we introduce the notation $\check{\theta}_{ij} = \check{Q}_{\hat{z}_1(i)\hat{z}_2(j)}$ with $\check{Q}_{ab} = \mathrm{sign}(\bar{Y}_{ab}(\hat{z}))\left(|\bar{Y}_{ab}(\hat{z})| \wedge M\right)$. Recall the definition of $\hat{Q}$ in Proposition A.1 with Y replaced by $\mathcal{Y}$. Then, we have

$$\|\hat{\theta} - \tilde{\theta}\|^2 \le 2\|\hat{\theta} - \check{\theta}\|^2 + 2\|\check{\theta} - \tilde{\theta}\|^2.$$

Since $\|\check{\theta} - \tilde{\theta}\|$ can be bounded by the exact argument in the proof of Lemma 5.1, it is sufficient to bound $\|\hat{\theta} - \check{\theta}\|^2$. By Jensen inequality,

$$\|\hat{\theta} - \check{\theta}\|^2 \le \sum_{ab} |\hat{z}^{-1}(a)||\hat{z}^{-1}(b)|(\bar{Y}_{ab}(\hat{z}) - \bar{\mathcal{Y}}_{ab}(\hat{z}))^2 \le \|Y - \mathcal{Y}\|^2.$$

Thus, the proof is complete. □

References

[1] E. M. Airoldi, T. B. Costa, and S. H. Chan. Stochastic blockmodel approximation of a graphon: Theory and consistent estimation. In *Advances in Neural Information Processing Systems*, pages 692–700, 2013.

[2] D. J. Aldous. Representations for partially exchangeable arrays of random variables. *Journal of Multivariate Analysis*, 11(4):581–598, 1981.

[3] J.-F. Cai, E. J. Candès, and Z. Shen. A singular value thresholding algorithm for matrix completion. *SIAM Journal on Optimization*, 20(4):1956–1982, 2010.

[4] E. J. Candes and Y. Plan. Matrix completion with noise. *Proceedings of the IEEE*, 98 (6):925–936, 2010.

[5] E. J. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(6):717–772, 2009.

[6] E. J. Candès and T. Tao. The power of convex relaxation: Near-optimal matrix completion. *Information Theory, IEEE Transactions on*, 56(5):2053–2080, 2010.

[7] D. Choi. Co-clustering of nonsmooth graphons. *arXiv preprint arXiv:1507.06352*, 2015.

[8] P. Diaconis and S. Janson. Graph limits and exchangeable random graphs. *arXiv preprint arXiv:0712.2749*, 2007.

[9] C. Gao, Y. Lu, and H. H. Zhou. Rate-optimal graphon estimation. *The Annals of Statistics*, 43(6):2624–2652, 2015.

[10] C. Gao, A. W. van der Vaart, and H. H. Zhou. A general framework for bayes structured linear models. *arXiv preprint arXiv:1506.02174*, 2015.

[11] A. Guntuboyina. Lower bounds for the minimax risk using f-divergences, and applications. *Information Theory, IEEE Transactions on*, 57(4):2386–2399, 2011.

[12] J. A. Hartigan. Direct clustering of a data matrix. *Journal of the American Statistical Association*, 67(337):123–129, 1972.

[13] P. W. Holland, K. B. Laskey, and S. Leinhardt. Stochastic blockmodels: First steps. *Social Networks*, 5(2):109–137, 1983.

[14] D. N. Hoover. Relations on probability spaces and arrays of random variables. *Preprint, Institute for Advanced Study, Princeton, NJ*, 2, 1979.

[15] O. Kallenberg. On the representation theorem for exchangeable arrays. *Journal of Multivariate Analysis*, 30(1):137–154, 1989.

[16] R. Keshavan, A. Montanari, and S. Oh. Matrix completion from noisy entries. In *Advances in Neural Information Processing Systems*, pages 952–960, 2009.

[17] R. Keshavan, A. Montanari, and S. Oh. Matrix completion from a few entries. *Information Theory, IEEE Transactions on*, 56(6):2980–2998, 2010.

[18] O. Klopp, A. B. Tsybakov, and N. Verzelen. Oracle inequalities for network models and sparse graphon estimation. *arXiv preprint arXiv:1507.04118*, 2015.

[19] V. Koltchinskii, K. Lounici, and A. B. Tsybakov. Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. *The Annals of Statistics*, 39(5):2302–2329, 2011.

[20] M. Lee, H. Shen, J. Z. Huang, and J. S. Marron. Biclustering via sparse singular value decomposition. *Biometrics*, 66(4):1087–1095, 2010.

[21] L. Lovász. *Large Networks and Graph Limits*, volume 60. American Mathematical Society, 2012.

[22] L. Lovász and B. Szegedy. Limits of dense graph sequences. *Journal of Combinatorial Theory, Series B*, 96(6):933–957, 2006.

[23] Y. Lu and H. H. Zhou. Minimax rates for estimating matrix products. *Preprint, Yale University*, 2015.

[24] Z. Ma and Y. Wu. Volume ratio, sparsity, and minimaxity under unitarily invariant norms. *Information Theory, IEEE Transactions on*, 61(12):6939–6956, 2015.

[25] P. Massart. *Concentration Inequalities and Model Selection*, volume 1896. Springer, 2007.

[26] S. C. Olhede and P. J. Wolfe. Network histograms and universality of blockmodel approximation. *Proceedings of the National Academy of Sciences*, 111(41):14722–14727, 2014.

[27] B. Recht. A simpler approach to matrix completion. *The Journal of Machine Learning Research*, 12:3413–3430, 2011.

[28] B. Recht, M. Fazel, and P. A. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Review*, 52(3):471–501, 2010.

[29] D. B. Rubin. Inference and missing data. *Biometrika*, 63(3):581–592, 1976.

[30] P. J. Wolfe and S. C. Olhede. Nonparametric graphon estimation. *arXiv preprint arXiv:1309.5936*, 2013.

[31] B. Yu. Assouad, Fano, and Le Cam. In *Festschrift for Lucien Le Cam*, pages 423–435. Springer, 1997.